

AD-A086 735

YALE UNIV NEW HAVEN CT DEPT OF COMPUTER SCIENCE
AFFECT UNITS AND NARRATIVE SUMMARIZATION.(U)
MAY 80 W 6 LEHNERT

F/6 5/7


N00014-75-C-1111

NL

UNCLASSIFIED RR-179

[06]
AD
AL-0735

END
DATE
FILMED
8-80
DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

ADA 086735

LEVEL

12



DTIC
ELECTE
JUL 14 1980
S D C

Affect Units and Narrative Summarization

Wendy G. Lehnert

Research Report #179

May 1980

This document has been approved
for public release and sale; its
distribution is unlimited.

DDC FILE COPY

YALE UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE

80 7 11 098

12

There is an error on page 23 of the report.
The last 2 lines of page 23 now read "CAT"; but should read "negative event".

SECRET

Affect Units and Narrative Summarization

Wendy G. Lehnert

Research Report #179

May 1980

This work was supported in part by the Advanced Research Project Agency under contract N00014-75-C-1111✓ and in part by the National Science Foundation under contract IST7918463.

This document has been approved
for public release and sale; its
distribution is unlimited.

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 1488-179	2. GOVT ACCESSION NO. AD-A086 735	3. RECIPIENT'S CATALOG NUMBER 14
4. TITLE (and Subtitle) Affect Units and Narrative Summarization		5. TYPE OF REPORT & PERIOD COVERED Research Report
7. AUTHOR(s) Wendy Lehnert		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Yale University-Computer Science Dept. PO Box 2158 - Yale Station New Haven, Ct 06520		8. CONTRACT OR GRANT NUMBER(s) N00014-75-C-1111 NSG-IST79-18463
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency Arlington, VA 22209		12. REPORT DATE May 86
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Program Arlington, VA 22217		13. NUMBER OF PAGES 53
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this report is unlimited		15. SECURITY CLASS. (of this Report) unclassified
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Natural Language Processing Memory Representation Summarization Narrative Plot Recognition		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The analysis of narrative text involves various levels of description. On the lowest level are word meanings and syntactic structure within single sentences. On a higher level there are problems of generating inferences and integrating information into memory. At the highest level is the notion of a macro-structure or narrative plot. The identification of high-level narrative structures		

DD FORM 1473

JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

-- OFFICIAL DISTRIBUTION LIST --

Defense Documentation Center Cameron Station Alexandria, Virginia 22314	12 copies
Office of Naval Research Information Systems Program Code 437 Arlington, Virginia 22217	2 copies
Dr. Judith Daly Advanced Research Projects Agency Cybernetics Technology Office 1400 Wilson Boulevard Arlington, Virginia 22209	3 copies
Office of Naval Research Branch Office - Boston 495 Summer Street Boston, Massachusetts 02210	1 copy
Office of Naval Research Branch Office - Chicago 536 South Clark Street Chicago, Illinois 60615	1 copy
Office of Naval Research Branch Office - Pasadena 1030 East Green Street Pasadena, California 91106	1 copy
Mr. Steven Wong New York Area Office 715 Broadway - 5th Floor New York, New York 10003	1 copy
Naval Research Laboratory Technical Information Division Code 2627 Washington, D.C. 20375	6 copies
Dr. A.L. Slafkosky Commandant of the Marine Corps Code RD-1 Washington, D.C. 20380	1 copy
Office of Naval Research Code 455 Arlington, Virginia 22217	1 copy

Office of Naval Research Code 458 Arlington, Virginia 22217	1 copy
Naval Electronics Laboratory Center Advanced Software Technology Division Code 5200 San Diego, California 92152	1 copy
Mr. E.H. Gleissner Naval Ship Research and Development Computation and Mathematics Department Bethesda, Maryland 20084	1 copy
Captain Grace M. Hopper NAICOM/MIS Planning Board Office of the Chief of Naval Operations Washington, D.C. 20350	1 copy
Dr. Robert Engelmores Advanced Research Project Agency Information Processing Techniques 1400 Wilson Boulevard Arlington, Virginia 22209	2 copies
Professor Omar Wing Columbia University in the City of New York Department of Electrical Engineering and Computer Science New York, New York 10027	1 copy
Office of Naval Research Assistant Chief for Technology Code 200 Arlington, Virginia 22217	1 copy
Major J.P. Pennell Headquarters, Marine Corp. (Attn: Code CCA-40) Washington, D.C. 20380	1 copy
Computer Systems Management, Inc. 1300 Wilson Boulevard, Suite 102 Arlington, Virginia 22209	5 copies
Ms. Robin Dillard Naval Ocean Systems Center C2 Information Processing Branch (Code 8242) 271 Catalina Boulevard San Diego, California 92152	1 copy

"Affect Units and Narrative Summarization"

Wendy G. Lehnert
Department of Computer Science
Yale University

ABSTRACT

The analysis of narrative text involves various levels of description. On the lowest level are word meanings and syntactic structure within single sentences. On a higher level there are problems of generating inferences and integrating information into memory. At the highest level is the notion of a macro-structure or narrative plot. The identification of high-level narrative structures is central to the problem of narrative summarization. But the intuitive notion of a plot is useless for a process model of summarization unless we can specify the hierarchical representations that allow us to analyze input and produce plots as output.

A representational strategy has been developed for high-level structural analysis in conjunction with the BORIS system (a narrative text understanding system). The structures produced effectively encode plot lines in terms of connected graph structures where graph nodes correspond to specific affect units. An affect unit is an abstraction of affective causality which is recognized in a bottom-up manner at the time of understanding. Simple manipulations of these graph structures yield the conceptual basis for narrative summaries, but the actual process of summarization depends on certain connectivity properties in the graph structure. Summarization techniques for the general case are presented, and a specific algorithm for one class of graphs is proposed.

1. INTRODUCTION

When a person reads a narrative story, an internal representation for that story is constructed in their memory. We can examine the contents of this memory representation by asking the reader simple questions about the story. Typical question answering behavior will reveal evidence for numerous inferences, causal chain constructions, and the predictive integration of information into instantiated knowledge structures [Lehnert 1978, Dyer and Lehnert 1980]. While question answering provides us with a method for examining the contents of a memory representation, the task of question answering does not readily yield a more global picture of the memory representation as a whole. We can only guess at how the various pieces fit together within a single structure.

If we are interested in the structure of narrative memory representations, we must turn to the task of summarization. When a reader is asked to summarize a story, vast amounts of information within the memory representation are selectively ignored in order to produce a distilled version of the original narrative. This process of simplification relies on a global structuring of memory that allows search procedures to concentrate on central elements of the story while ignoring peripheral details. We intuitively expect that some hierarchical structure is holding memory together, but the precise formulation of this structure is much more elusive. Any process model

The author is indebted to Chris Riesbeck and Roger Schank for their comments on a rough draft of this paper.

that attempts to utilize high level narrative structures must confront a number of difficult questions. How is the hierarchical ordering of a memory representation constructed at the time of understanding? Exactly what elements of the memory representation are critical in building this structure? What search processes are used to examine memory during summarization? How are summaries produced after memory has been accessed? In this paper we will propose a method for narrative analysis and summarization that addresses each of these issues.

2. MOTIVATION AND BACKGROUND

A theoretical tradition of language processing efforts is represented by a succession of language processing programs known as the MARGIE [Schank 1975], SAM [Cullingford 1977], and PAM [Wilensky 1978] systems. The most recent effort in this tradition is the BORIS system [Dyer and Lehnert 1980] which processes narratives using multiple knowledge structures (scripts, plans, goals, and themes) as outlined in [Schank & Abelson 1977].

In addition to the knowledge structures proposed by Schank and Abelson, we have attempted to implement a theory of affect inference [Roseman 1979] so BORIS can recognize and predict emotional reactions on the part of its narrative characters. Roseman's system is based on a fairly simple thesis: emotional reactions to an event can be predicted by decomposing the event along five dimensions: (1) desirability, (2) attainment, (3) agency, (4) legitimacy, and (5) certainty. This system of decomposition suggested a representational strategy for events that could be partially implemented in BORIS. But

the representation needed for our purposes is even simpler than the five-way decomposition of Roseman. A high level analysis of narrative structures can be built on the basis of distinctions in only the first two dimensions: desirability and attainment.

In the next four sections we will develop a representational system for high level structural analysis. We will then examine the conceptual content of narrative summaries and present a rigorous framework for describing narrative cohesion. Once this preparation is complete, we can propose a process-oriented description of summary generation based on structural analysis. We will then complete the process model by describing recognition techniques needed to produce high level structural analyses. The conclusion will present a brief comparison between the proposed system and other summarization strategies, and close with some comments on other applications for high level structural analysis.

2. PRIMARY AFFECT STATES

We can reduce the notion of a spectrum into a binary value by dividing the spectrum in half and choosing sides. An event is either desirable or undesirable, and its resulting states have either been attained or not. Using this reduction, the two dimensions of desirability and attainment produce four states:

		DESIRABILITY	
		+	-
ATTAINMENT	+	GDT	GAT
	-	LDT	LAT

These states correspond to four intuitive notions:

GDT (Getting Desired Thing) Events that satisfy desire
 GAT (Getting Aversive Thing) Events that create problems
 LDT (Lacking Desired Thing) Internal states of desire
 LAT (Lacking Aversive Thing) Internal states of satisfaction

For example, an initiated goal is an LDT while a satisfied goal is a GDT. The termination of a positive interpersonal theme constitutes a GAT, and the termination of a negative interpersonal theme yields an LAT.


Initially, all four states were used as the building blocks for high level narrative structures. But after numerous sample texts had been analyzed and our representational needs began to solidify, it became apparent that the identification of LAT's was superfluous. The initial set of four states therefore gave way to a representational system based on three affect states. We will rename these states in a suitably suggestive manner:

OO (Positive Event) Events that please
 XX (Negative Event) Events that displease
 MM (Mental State) Internal states of desire

For the purposes of affect analysis, it was necessary to create linear maps of affect states for each character as the story progressed. After examining a number of these affect maps, patterns began to emerge of powerful generality. For example, consider the affect map for John in the following story:


When John tried to start his car this morning, it wouldn't turn over. He asked his neighbor Paul for help. Paul did something to the carburetor and got it going. John thanked Paul and drove to work.

The affect analysis for John consists of three affect states:




These three affect states are connected by a structure of three pairwise links. An aversive event motivated John to get his car started, John achieved this desire by getting Paul to start the car, and Paul's assistance resolved the original difficulty. This pattern represents an affect configuration that is extremely pervasive in narrative texts: resolution of a problem by intentional means.


Now suppose we extend our analysis to include the affect states of Paul:



The diagonal links signify causalities of affect across characters. When Paul agrees to help he is assuming John's state of desire; now Paul wants to get the car started too. This configuration of embedded achievement across characters signals cooperative service. When this configuration is followed by its symmetrical counterpart, we have an instance of exchange. Favors (voluntary services) are often exchanged, and the idea of a loan is a special case of exchanged services:



If we were missing the last two positive events in this structure, we would understand that A is obligated to B; it would be up to A to complete the symmetry of the configuration. A trade is also a special case of an exchange, where the two transactions occur simultaneously:



A number of standard affect configurations arise in this manner which allow us to recognize narrative structures and build plot lines from affect states. But before we can identify standard configurations, we must present a system of affect links that will be used to join pairs of affect states.

3. AFFECT LINKS

A link which runs from a negative event to a mental state describes motivation, while a link running from a mental state to a positive event describes actualization. To make these and other distinctions explicit, we will use a system of four link types: MOTIVATION (m), ACTUALIZATION (a), TERMINATION (t), and EQUIVALENCE (e). Each link describes an oriented arc between two affect states. M-links and a-links point forward in time, while t-links and e-links point backward in time. With three affect states and four link types, there are 36 possible pairwise configurations if we consider all the possible combinations. But in fact, only 15 of these will occur when we observe some syntactic constraints on the use of affect links. To summarize these constraints, the following table illustrates which combinations occur by marking legal configurations with an "*".

	(M) (M)	(M) (O)	(M) (X)	(O) (M)	(X) (M)	(O) (X)	(X) (O)	(O) (O)	(X) (X)
m	*			*	*				
a		*	*						
t	*					*	*	*	*
e	*					*	*	*	*

The constraints can be described as follows:

- m-links must point to a mental state.
- a-links must point from a mental state to an event.
- t-links and e-links must point:
 - a) from a mental state to a mental state, or
 - b) from an event to an event.

These links have been given an orientation for intuitive convenience rather than notational necessity. For m-links and a-links, the pointer moves from a temporal antecedent to its consequent. With t-links, the pointer goes from a subsequent event to the prior event it terminates. E-links are used to identify redundant state descriptions which appear at different times. The backward orientation of e-links is therefore arbitrary.

4. PRIMITIVE AFFECT UNITS

Our 15 legal pairwise configurations will act as the building blocks for more complex configurations. We will refer to them as the "primitive affect units," and each will be specified by name:

SUBGOAL	SUCCESS	FAILURE
MM \curvearrowright m MM \curvearrowright m	MM \curvearrowright a OO \curvearrowright a	MM \curvearrowright a XX \curvearrowright a
CHANGE OF MIND	LOSS	MIXED BLESSING
MM \curvearrowright t MM \curvearrowright t	OO \curvearrowright t XX \curvearrowright t	OO \curvearrowright e XX \curvearrowright e

PERSEVERENCE

MM \int_e
MM

RESOLUTION

XX \int_t
00

HIDDEN BLESSING

XX \int_e
00

ENABLEMENT

00 \int_m
MM

NEG. TRADE-OFF

00 \int_t
00

COMPLEX POS. EVENT

00 \int_e
00

MOTIVATION

XX \int_m
MM

POS. TRADE-OFF

XX \int_t
XX

COMPLEX NEG. EVENT

XX \int_e
XX

Sometimes a primitive affect unit will appear without other interceding affect states. This occurs most commonly with the units of motivation, enablement, and subgoals. The other primitive affect units tend to be broken up by interceding affect states. For example, it may take months (with lots of interceding emotional reactions) to find out that a job promotion is now leading to an ulcer. This would be an example of a mixed blessing, a good thing turned sour.

EXAMPLES OF PRIMITIVE AFFECT UNITS

MOTIVATION: You get fired and need a job.
You bounce a check and need to deposit funds.
Your wife dies and you long for companionship.

SUCCESS: You decide to ask for a raise and you get it.
You fix a flat tire.
You need a car so you steal one.


FAILURE: Your proposal of marriage is declined.
You can't find your wallet.
You can't get a bank loan.

RESOLUTION: Your broken radio starts working again.
They catch a thief who got your wallet.
You fix a flat tire after a blow out.


LOSS: Your big income tax refund is a mistake.
Your wife gets a divorce.
The car you just bought is totaled.

5. COMPLEX AFFECT UNITS

Using the 15 primitive affect units, we can build larger affect units to represent general plot configurations. For example, the string (X M O) of 3 affect states is used by 3 different affect units that are distinguished only by the links involved:

INTENTIONAL
PROBLEM RESOLUTION

= motivation
& success
& resolution

FORTUITOUS
PROBLEM RESOLUTION

= motivation
& resolution

SUCCESS BORN
OF ADVERSITY

= motivation
& success


These are examples of complex affect units that are commonly found at the center of plot structures. Other closely related affect units include:

FLEETING SUCCESS



= success
& loss

STARTING OVER



= success
& loss
& motivation
& perseverance


GIVING UP



= failure
& motivation
& change of mind


Many complex affect units can be transformed into different units by way of a minor variation:

SACRIFICE




= success
& trade-off

NESTED SUBGOALS



= subgoal
& success
& success

KILLING TWO BIRDS



= complex pos. event
& success
& success


Thus far we have concentrated on affect units that describe configurations within a single character. A large number of complex affect units involve multiple characters. Affect units with more than one character require cross-character affect links. These will be represented by diagonal segments between affect states, where the higher affect state precedes the lower affect state in time. While we found it useful to distinguish four types of intra-character links in building the primitive affect states, we will not need to distinguish cross-character links. Cross character links can occur between any pair of affect states, and their interpretation will rely on the following conventions:

RESULTING MENTAL STATES:


REQUEST



ENABLEMENT



MOTIVATION



These configurations describe the initiation of a goal state as a direct response to another character's situation. All of the resulting mental states are initiated by free choice. In the case of "MM/MM", the resulting mental state occurs in response to a request. This resulting mental state may assume the desires of the initiator, or it may oppose them. The request configuration does not commit us

to any assumptions about the contents of the two mental states or how their contents are related. In the cases of "++/MM" and "--/MM", we have mental states enabled or motivated by events. For example, a desire to celebrate is normally enabled by a positive event, while a desire to help out is typically motivated by a negative state.

SPEECH ACTS




These two configurations describe communications which result in positive and negative affect states. The antecedent in either case is a mental state describing the intentions of that character. These two configurations often appear in tandem when an agreement is achieved by coercion, i.e., a promise is motivated by a threat.




Shared events are shared in the sense that two characters are affected by them in a similar manner. The same event is experienced by both people as either a positive event or a negative event. Mixed events are just the opposite. Here the same event is experienced differently by both people, one is affected positively, and one negatively.

These nine cross-character configurations can now be used to build complex affect units involving two characters. Some of the most common configurations involving two characters are those that describe cooperative agreements and behavior. In the simplest case a request is made and the respondent behaves either cooperatively or not:


HONORED REQUEST



DENIED REQUEST




BUNGLED REQUEST




In our story about John's car we had an instance of an honored request. John asked Paul for help and Paul got the car started for John. When this situation occurs we can assume that the second character assumes the mental goal state of the first character. Paul wanted to get the car started too. When a request is denied, we should assume different mental states. If Paul tells John that he's too busy, we should not assume that Paul wanted to get John's car started.

A slightly different situation arises when the speech act of a threat is invoked instead of a request:


EFFECTIVE COERCION



INEFFECTIVE COERCION




BUNGLED COERCION




In these situations the respondent is confronted with a problem situation that can be resolved with either cooperative behavior or a challenging denial. These situations are very common, and in some cases it is appropriate to represent them in greater detail. For example, what if Paul agrees to help John get his car started, but then fails to do so? In some stories, the extraction of an agreement receives enough attention to warrant its own affect analysis:

PROMISED REQUEST HONORED




PROMISED REQUEST BUNGLED




When John asks Paul for help he has set up a subgoal for getting help. If Paul agrees to help, Paul satisfies John's subgoal by making a promise. If Paul then succeeds in helping John, the top level goal is achieved as well. But if Paul fails, his actions amount to nothing more than good intentions that were bungled. The affect units for an honored request and a promised request that is honored are very similar. When a request is honored, we have a request and shared success. When a request is promised and then honored, we have nested subgoals, a request, a promise, and shared success. These are identical except for details about the agreement as an interaction that is separate from the service performed. This more detailed level of description is necessary when we try to represent "good intentions" that fail in response to a request or threat.

If we examine the notion of a threat at this level of detail we can see the difference between a threat that is agreed to and successful, versus a threat that is sincerely agreed to but unsuccessful anyway. (These are elaborations on effective coercion and bungled coercion).

COERCED AGREEMENT HONORED




COERCED AGREEMENT BUNGLED




In both of these cases the respondent intends to go along with the threat. When the threat succeeds, it is because the respondent succeeds. When the threat fails, it is because the respondent fails. In both cases the respondent promises to cooperate. A slightly different situation arises when the respondent promises to cooperate but intentionally fails to come through:

DOUBLE-CROSS




In a double-cross, the respondent deceptively agrees to go along, and then intentionally does something to foil the other's goal. This unit contains subgoals, a request, a promise, and a mixed event of success and failure. We could also represent a double-cross in response to coercion if the request were replaced with a threat:

COERCED DOUBLE-CROSS



This version of a double-cross seems more defensible since it was initiated by a coercive act. We can see symmetry in the negative consequences to both characters. In addition to cooperative and uncooperative responses, people often interact in unsolicited ways:


UNSOLICITED HELP



In this case the problem state is completely assumed by an intervening character who is motivated by the initial problem state to initiate his own assistance. If Paul had noticed John's problem and volunteered his services we would have a case of unsolicited help.


Any of the preceding affect units for cooperative behavior can be embedded in a problem resolution. For example, a problem resolution via a successful threat would look like:

PROBLEM RESOLUTION BY EFFECTIVE COERCION




In addition to the various ways that one character can react to another's desires, there are also a number of standard affect configurations that describe situations of reciprocation. When cooperative behavior is reciprocated, we arrive at affect units for obligation, exchange, and trades:


OBLIGATION



SERIAL EXCHANGE



SIMULTANEOUS EXCHANGE




Notice that the unit for a serial exchange (of requests) is very similar to the unit for a simultaneous exchange (or trade). In a serial exchange the requests are satisfied one after another, while in


a simultaneous exchange requests are handled in parallel. The same affect states and link configurations occur in both units; only the temporal sequencing of the affect states is different.

Another variation on exchanged requests occurs when the respondent agrees to honor the initial request pending a conditional request of his own. Paul could have agreed to fix John's car if John would first give him a beer. Then we would have two requests with one being conditional on the completion of the other:

REQUEST HONORED WITH
CONDITIONAL REQUEST



REQUEST HONORED WITH
CONDITIONAL PROMISE




Of course John may only promise to give Paul a beer. In this case the request is met with a conditional promise. While we expect John to honor his promise, he may not. If he doesn't, we will find the pattern for a double-cross.


These units for conditional requests illustrate a point about the "syntax" of affect configurations. Whenever a subgoal unit occurs, the subgoal must be satisfied before the supergoal. This normally happens with two success units, yielding the complex unit for nested subgoals (as happens in a request honored with a conditional request). But occasionally a promise will suffice for subgoal satisfaction (as is the case for a request honored with a conditional promise). So subgoals can be satisfied in one of two ways, by a success unit or a

promise unit. These two methods of satisfaction are nevertheless distinguished by predictive processing. Whenever a promise unit is encountered, we must activate expectations for success resulting in a positive shared event, or success resulting in a positive mixed event. That is, a primitive unit for a promise always sets up expectations for the complex units describing an honored promise or a reneged promise. In either case, the promise unit and the success unit would share the same mental state (driving both the promise and its actualization).

HONORED PROMISE




RENEGED PROMISE




Promises and cooperative behavior are not the only affect units that rely on shared and mixed events. Other complex affect units include:


MALICIOUS ACT



KIND ACT




COMPETITION




Finally, a number of affect units involve variable affect states:


RETALIATION



REGRETTABLE MISTAKE



SABOTAGE



The unspecified affect state here signifies a "wild card" for the purposes of pattern recognition. Any affect state will match an unspecified state.

This section has attempted to show how complex units can be constructed to provide infinite variations of plot lines. For example, a kind act with a resulting trade-off will amount to an act of self-sacrifice, while a fortuitous problem resolution with a trade-off will merely signify an undesirable side-effect. It would be pointless to try to enumerate at this time all of the possible combinations that may be useful for structural recognition.


6. SUMMARIZATION

By recognizing affect units, we can achieve a high level analysis of activities and interactions within a narrative. We should expect to find evidence for this "chunking" of information in paraphrase and summarization behavior. To see how this works, consider the following narrative:

John was thrilled when Mary accepted his engagement ring. But when he found out about her father's illegal mail-order business, he felt torn between his love for Mary and his responsibility as a policeman. When John finally arrested the old man, Mary called off the engagement.

M loves J
engagement

arrest made
wants revenge
engagement off



J loves M
engagement
discovers crook
wants to enforce law
makes an arrest

engagement off

The affect analysis for John and Mary reveal configurations of a trade-off from retaliation on the part of Mary, and a problem resolution leading to loss for John.* Ideally, one might expect good summaries to convey each of these four affect units (trade-off, retaliation, problem resolution, and loss). A stronger claim about summaries would argue that any summary which does not convey all four affect units is an inferior summary.

For example:

"When John arrested Mary's father, she interfered with his wedding."
(no trade-off for Mary)

"When John arrested an old crook, Mary called off their engagement."
(no retaliation for Mary)

"When Mary's father was arrested, she called off her engagement."
(no problem resolution for John)

"When John arrested Mary's father, she called off her engagement."
(no loss for John)

But a summary that includes all four affect units provides an accurate description of the story:


"When John arrested Mary's father, she called off their engagement."
(all units present)

Of course "inclusion" here means inclusion by inference as well as by explicit mention. We must infer that there is a causality between John's act and Mary's act in order to understand retaliation, but this inference had to be made with the original narrative as well.

* We will ignore the initial success units for reasons that will be explained in section 9.

Ultimately, an affect analysis in terms of affect units should allow us to predict the sorts of summaries that human subjects will produce. But initially, we must study actual summary behavior in order to develop a process model that converts affect unit configurations into narrative summaries. Consider the following story:

John and Bill were competing for the same job promotion at IBM. John got the promotion and Bill decided to leave IBM to start his own consulting firm, COMSYS. Within three years COMSYS was flourishing. By that time John had become dissatisfied with IBM so he asked Bill for a job. Bill spitefully turned him down.



Here we have a competitive situation between John and Bill in which John wins. Bill's failure turns into success out of adversity, and then he retaliates against John for his initial failure. John sets the stage for Bill's retaliation by asking Bill for a job. John consequently experiences a failure when Bill uses this opportunity to get revenge by denying John's request. The affect units here are (1) competition which subsumes (2) John's success and (3) Bill's failure, (4) success born of failure, (5) retaliation, and (6) a denied request which subsumes (7) John's request and (8) Bill's denial. To see how these units are integrated into summaries of the story, we will look at 10 summaries provided by experimental subjects. The subjects were asked to read the story, and were then instructed to summarize the story in one sentence. The summaries which appear below are verbatim

responses, except for some name corrections (subjects frequently reversed Bill and John).

SUMMARIZATION BEHAVIOR

- (1) John, Bill compete for a job which John wins, causing Bill to quit the company and start his own firm (COMSYS), which leads to Bill's spiteful rejection of John's request for a job some years later at Bill's successful company.
- (2) Bill was spiteful when John asked him for a job, because they had once competed for the same job at IBM.
- (3) John and Bill were both competing for a job at IBM which John got so Bill started his own business and later had the opportunity to turn John down when John wanted a job.
- (4) John got promoted by IBM, so Bill his friend, started his own business which soon flourished and when John came asking for a job, Bill spitefully turned him down.
- (5) Bill and John worked for IBM, and were friends until 3 years later Bill turned John down when he asked for a job.
- (6) Bill turned John down for a job because John had beat him out of a promotion when they both worked for IBM.
- (7) Bill started his own business COMSYS after losing out to John for a job at IBM and later out of spite refused to give John a job when John was dissatisfied with his old one.
- (8) Bill, who lost a job promotion to his competitor John, establishes a lucrative consulting firm of his own, and rejects John's request for a job later on.
- (9) John beat out Bill for a promotion at IBM whereupon Bill decided to leave and form his own company, COMSYS, which was flourishing within 3 years, and which John turned to for a job when he was fed up at IBM which he did not get due to Bill's spite.
- (10) John was promoted at IBM instead of Bill, so when Bill left IBM to start his own firm and the business flourished, he turned John down when the latter, dissatisfied at IBM, applied to Bill for a job.

In analyzing these summaries for the presence of affect units, we find that a number of units are present only in an implicit manner. For example, (1,2,4,7,9) explicitly refer to "spite" and therefore make explicit reference to the retaliation unit. Summaries (3,6,8,10) are constructed with suggestive causalities, and the presence of retaliation is only implicitly present. These implicit cases can be contrasted with (5) where there is no basis for a retaliation unit whatsoever. When retaliation is implicit, it is conveyed by the causal constructions of clause formation. Other affect units may be implicitly present by conceptual entailment. For example, in (3, 6, 7) the request (for a job) is implicit from the verb phrase "to turn down," since this expression describes a denied request. In all of the other summaries John's request is explicit. In the chart on the next page we have marked with an "IMP" those affect units which are implicit in the text.

Other affect units are present by processes of inference. Examples of inferred affect units occur in (2) and are marked with an "INF." The explicit presence of retaliation and competition force us to infer that John won and Bill lost. The patterns of competition and retaliation wouldn't overlap at a negative event if Bill got the job. Without this overlap, we would say that it just doesn't make sense for Bill to get the job and then feel spiteful about it. This inference is a "role-binding inference," driven by the retaliation unit.

"X is spiteful toward Y" sets us up for:

- (1) a causal antecedent: Y causes a GAT for X, and
- (2) a causal consequent: X causes a GAT for Y.

Since competitive resolution entails the configuration needed by (1),



we can establish who is the winner and who is loser by a role binding inference (If X is spiteful toward Y, X is the loser and Y is the winner). So we can say that Bill's failure and John's success are present by implicit inference in summary (2). The presence of denial is also supplied by the retaliation unit where the structure for a negative mixed event is encoded.

If we analyze these summaries for the presence of our 8 affect units, we get the following distribution:

	competition (COMP)				retaliation (RET)			
	Bill's success (BS)				John's success (JS)			
	Bill's failure (BF)				denied request (DR)			
	John's failure (JF)				John's request (JR)			
	COMP	BS	RET	DR	BF	JS	JR	JF
1	X	X	X	X	X	X	X	X
2	X		X	INF	INF	INF	X	INF
3	X	X	IMP	X	X	X	IMP	X
4		X	X	X		X	X	X
5				X			X	X
6	X		IMP	X	X	X	IMP	X
7	X	X	X	X	X	X	IMP	X
8	X	X	IMP	X	X	X	X	X
9	X	X	X	INF	X	X	X	X
10	X	X	IMP	X	X	X	X	X

We could postulate a rough qualitative ranking of the summaries based on the number of affect units present. The number of summaries conveying affect units breaks down as follows:

summaries	affect units
6	8
2	7
1	6
1	3

This distribution suggests that the affect unit analysis is central to paraphrase production. Summary (5) with three affect units does seem to be the poorest summary, while the others (while stylistically different) are more on par in terms of their content.

By analyzing the nature of affect unit occurrences in terms of their explicit expression or implicit presence, we begin to see that some units are "pivotal" in driving inferences about other units. The identification of pivotal units will be very important in the actual process of summarization. We will return to this idea in section 9 when we outline the process model for summarization.

To see how summaries are built from affect configurations, we can look at these 10 summaries in terms of their clause constructions. In the following abstractions, we have abbreviated all clauses that describe affect units and identified them accordingly. What remains is a structural backbone for the sentences generated:

- (1) [COMP] which [JS,BF] causing [BS] which leads to [RET,DR,JF] of [JR].
- (2) [RET] when [JR,] because [COMP]. (infer JS,BF,DR,JF)
- (3) [COMP] which [JS,BF] so [BS] and later had the opportunity to [DR,JF]. (implicit JR,RET)
- (4) [JS] so [BS] and when [JR], [DR,RET,JF].
- (5) [DR,JF] when [JR].
- (6) [DR,JF] because [COMP,JS,BF]. (implicit JR,RET)
- (7) [BS] after [COMP,JS,BF] and later [DR,RET,JF]. (implicit JR)


- (8) Bill, who [COMP,JS,BF], [BS] and [DR,JF,JR] later on. (implicit RET)
- (9) [COMP,BF,JS] whereupon [BS] and which [JR] which [JF] due to [RET]. (infer DR)
- (10) [COMP,BF,JS] so when [BS], [DR,JF] when [JR]. (implicit RET)

These skeletons reveal natural "clumps" of information. For example, Bill's failure and John's success are naturally tied to their competition. This follows from the fact that competition entails units for success and failure. If Bill's success is mentioned at all, it occurs in isolation of other units, and always follows the COMP-JS-BF clump (when their order is inverted in (7) the connective makes their relationship explicit). John's failure and his denied request tend to appear together, and can be easily combined with retaliation when retaliation is made explicit. The choice of specific connectors appears to be determined by retaliation, since the causality connecting other affect units serves to convey retaliation implicitly. The interplay between global factors (like retaliation) and more local entities (like John's job request) can be handled in a variety of ways. Some constructions are stylistically more pleasing than others, and the use of implicit and inferential information seems central to the more successful strategies.

7. NARRATIVE COHESION

It is possible to assess the cohesiveness of a narrative by analyzing its connectivity across affect units. For example, in the COMSYS story, we have a totally coherent text: John's success causes Bill's failure and this motivates Bill to become successful on his own. Bill then exploits an opportunity to retaliate against John by


causing John to fail in his job hunting. The causal chain is not quite linear, but it is completely connected:



Suppose the last sentence of the COMSYS story was:

"Bill gave John a key position in his company."

Then we would have a slightly different set of affect units:



Now we no longer have retaliation. Instead, we have Bill honoring a request by John which yields John another success. The story is no less plausible, but its cohesiveness is lessened; now there is no connectivity between the first three and the last two affect units. Bill's success enables him to help John, but there is no affect-oriented connection to unify the story.



We would tend to say that Bill gave John a good job in spite of the fact that John won the promotion Bill wanted. We are more surprised to see Bill act magnanimously; a retaliation seems more likely. But


this expectation is not founded on any knowledge of Bill's personality or attitude toward John. We have no such information to help us predict his behavior. It is instead a general expectation for narrative unity. We have a preference for cohesive narratives, and retaliation allows us to tie everything together. If Bill offers John the job, we cannot establish total connectivity across all affect units.

This type of expectation derives from our knowledge about narratives rather than our knowledge about the world in general. It is a weak expectation in the sense that it can be easily overridden by specific knowledge. For example, if we knew that Bill and John were best friends throughout the story, then we would not expect retaliation.

It could be argued that any expectation for retaliation in the COMSYS story is really an expectation about an eye for an eye rather than narrative unity. Bill was John's rival and Bill will want to get even. This level of expectation relates to the symmetry of a story. A story is weakly cohesive if it has a symmetry in its cross-character affect causalities. Retaliation is an affect unit that completes the symmetry of aversive causalities. When Bill refuses John the job, we have both strong cohesion (total connectivity across affect units) and weak cohesion (symmetry in the cross-character affect links). But when Bill offers John a good job, we have neither strong nor weak cohesion.

When a narrative embodies total symmetry, we detect this immediately and remember it as a salient feature of the story. For example, consider "The Gift of the Magi," by O. Henry. This is a story about a young couple who want to buy each other Christmas presents. They are both very poor. Della has long beautiful hair, and Jim has a prized pocket watch. To get money for the presents, Della sells her hair and Jim sells his pocket watch. Then she buys him a gold chain for his watch, and he buys her an expensive ornament for her hair. When they find out what they've done, they are consoled by the love behind each other's sacrifices.

This story exhibits an extreme symmetry:



This configuration involves (1) nested subgoals and (2) achievement (in getting and giving the gifts), (3) loss (in no longer having the things they sold), (4) another loss (in no longer having pleasure from the act of giving) (5) regrettable mistakes (the bad gifts), and (6) hidden blessings (in realizing what the gifts signify). Not only is there complete symmetry across both characters, but there are ironic causalities across the affect units. For example, the sense of loss does not occur until the top-level goals are achieved (when the gifts are exchanged). At the same time, this loss is also the basis for a

hidden blessing at the end of the story when they realize how the gifts signify their unselfish love for each other.

Symmetries of this sort can be heavily exploited in long term memory representations. For example, a long term recall of this story might include the fact that (1) Della sold her hair to buy Jim a gift, and (2) Jim bought Della an ornament for her hair. If these facts are augmented by knowledge of symmetry, a subject might then remember that (3) Jim sold X in order to buy the ornament, and (4) Della's gift to Jim was no longer appropriate after he sold X. If (3) and (4) were remembered by symmetric reconstruction, the actual identity of X and Della's gift might be forgotten.

Narrative cohesion will be an important factor for effective memory retention: cohesive texts (as defined by connectivity across affect units) should be remembered with greater accuracy than non-cohesive texts. While this claim is not central to the problem of text summarization, we can expect the two problems to be strongly related. But before we can proceed with either problem, we must become a bit more rigorous about the notion of connectivity across affect units.

8. CONNECTIVITY DEFINED

This section will develop the terminology necessary for a precise statement of our process model. As humans, we can look at graphic affect representations for narratives, and perceive rough degrees of connectivity within those representations. But a computational model that relies on connectivity will have to manipulate a precise

formulation of connectivity. So we must now resort to a few dry definitions. Once we have a suitably precise terminology, the actual process model will follow with relative ease.

In all that follows, let A and B be affect units.

DEFINITION: A is related to B if and only if A and B share a common affect state. (for convenience, we assume $A = B$)

DEFINITION: A is connected to B if and only if one of the following conditions hold:

- a) $A = B$
- b) A is related to B
- c) there is a sequence of affect units U_1, \dots, U_n such that A is related to U_1 , U_i is related to U_{i+1} , and U_n is related to B.

DEFINITION: A family around A is the set of affect units that are related to A. The family around A will be designated as $F(A)$.

DEFINITION: A cluster around A is the set of affect units that are connected to A.

In all that follows, let F be a family and K be a cluster.

DEFINITION: A entails B if and only if all affect states contained in B are also contained in A. (we may say that A entails B or that B is entailed by A).

DEFINITION: Let A be an affect unit contained in K. A is a top level affect unit in K if and only if A is not entailed by any other affect units contained in K.

DEFINITION: The size of K is the number of top level affect units contained in K. The size of K will be designated as $o(K)$.

DEFINITION: F is a maximal family in K if and only if F is a family contained in K and $o(F) \geq o(G)$ for all families G contained in K.

DEFINITION: A is a pivotal unit in K if and only if the family around A is a maximal family in K.

DEFINITION: K is a simple cluster if and only if K has one pivotal unit.

DEFINITION: We will define a metric on K as follows:

- (i) $d(A,B) = 0$ iff $A = B$
- (ii) $d(A,B) = 1$ iff A is related to B
- (iii) $d(A,B) = k$ iff U_1, \dots, U_{k-1} is the shortest sequence of affect units connecting A and B.

DEFINITION: The span of K is defined as the $\max\{d(A,B) \mid A \text{ and } B \text{ are units in } K\}$


DEFINITION: Let K be a simple cluster. The depth of K is defined as the $\max\{d(A,B) \mid A \text{ is the pivotal unit, } B \text{ is a unit in } K\}$.

These definitions describe simple graph structures that can be readily recognized in pictorial representations. We will look at three examples of affect unit graphs, but first, a few observations:

- (1) Maximal families may contain more than one pivotal unit. It is therefore possible to have a cluster with a unique maximal family that is not a simple cluster.
- (2) The definition for relatedness describes the simplest condition possible. We may later need to refine this to distinguish units that share n affect states ($n = 1, 2, 3$, etc.) and units whose shared affect states have certain properties of connectivity in terms of the affect links between them.
- (3) The notion of a top level unit is relative to the specification of some set of affect units. This allows us to examine the effect that different set specifications have on summarization behavior. For example, if we didn't include a unit for a denied request, the top level units for that configuration would drop down to the request, success, and positive mixed event. Various set specifications might be a key to individual differences in summary behavior.
- (4) More entailment between affect units results in simpler graph structures.
- (5) Larger units (in terms of affect states) are likely to result in greater connectivity as well as simpler graph structures.

The best way to get a sense of all this is to play with some concrete examples of the definitions in action.

Consider the story of John's broken engagement:




There are six top level affect units:

Mary's success	[MS]
John's success	[JS]
Mary's trade-off	[TO]
Mary's retaliation	[RET]
John's resolution	[RES]
John's loss	[JL]


The families for these units can be represented with a connectivity graph:

$F(MS) = \{TO\}$
 $F(JS) = \{JL\}$
 $F(RES) = \{RET\}$
 $F(JL) = \{RET, JS\}$
 $F(TO) = \{RET, MS\}$
 $F(RET) = \{TO, RES, JL\}$



There is only one maximal family and one pivotal unit (RET). This makes the cluster of 6 units a simple cluster. It has a depth of 2 and a span of 4.

Now consider the COMSYS story.




There are four top level affect units:

success born of failure	[SBF]
competition	[COM]
retaliation	[RET]
request denied	[RD]


The connectivity graph for these families is:

$F\{RD\} = \{RET\}$
 $F\{COM\} = \{RET, SBF\}$
 $F\{SBF\} = \{COM, RET\}$
 $F\{RET\} = \{COM, RD, SBF\}$



There is only one maximal family and one pivotal unit (RET). This yields a simple cluster with a depth of 1 and a span of 2.


Both of the previous stories result in fairly simple affect connectivity. For our last example, we will look at "The Gift of the Magi."



Now we have 10 top level affect units. Because of the symmetry of the story, we see the same units appearing for both the husband and the wife. These will be prefaced with "H" and "W" to signify which is which.

[HN]	nested subgoals	[WN]
[HM]	regrettable mistake	[WM]
[HL1]	loss of object	[WL1]
[HL2]	loss of achievement	[WL2]
[HB]	hidden blessing	[WB]

The connectivity graph for these looks like:



There are two maximal families, $F(HM)$ and $F(WM)$, with pivotal units HM and WM . Because this is not a simple cluster, the depth of the cluster is not defined. The span is 5.

9. SUMMARY GENERATION

Once the connectivity of the affect units has been established, we can drive a process of summary generation based on affect analysis. The generation of summaries for arbitrary narratives will require an extensive process model that can handle various classes of affect unit configurations. For example, the summary process for a simple cluster will have to be different from the summary process for a cluster with multiple pivotal units.

We will not attempt to present the complete solution here. But we can discuss the solution for a simple case in order to illustrate the techniques needed for the general case. We will therefore outline the process model for summaries of simple clusters, and then discuss methods for extending this solution to arbitrary clusters.

9.1 SUMMARIZATION OF SIMPLE CLUSTERS

The algorithm for generating summaries of simple clusters is a five step process:

- STEP 1: Find all top level affect units in K.
- STEP 2: Derive the affect unit graph structure.
- STEP 3: Identify the pivotal unit, P.
- STEP 4: Generate a base-line summary (S) from a frame for P.
- STEP 5: Augment S with information from affect units related to P.

Steps 1, 2, and 3 are simple manipulations based on the definitions of section 6.* Steps 4 and 5 require some explanation.


*(Step 3 involves summing the rows and columns of an adjacency matrix)

9.1.1 AFFECT UNIT FRAMES

All affect units are associated with generational frames (AUF's). These frames essentially designate how the affect unit can be expressed in natural language. For example, a frame for the "competition unit" might look like:

X and Y both [MM1], but Y [00].

where X, Y, MM1, and 00 are slots in the affect configuration:




This frame would give us summaries like, "Fred and Hank both loved Mary, but Hank married her." Or, "Bill and John both wanted the same job at IBM, but John got it." While this is a general frame that can be applied to any situation of competition, this frame may be overridden by knowledge-specific frames which are dependent on the specific instantiations of the our affect states. For example, when the competition is for a job promotion, we can say "Y was promoted over X at IBM" and this will convey the entire competition unit as well. This knowledge-specific frame can be invoked whenever the concept of a promotion appears in the parallel LDT's by a scheme of double indices on competition and promotion. The specification and selection of knowledge-specific generational frames will be a major problem for the production of smooth summaries.

9.1.2 INTEGRATING RELATED UNITS

Once a generational frame is chosen for the pivotal unit, we will transform the resulting base-line summary into a final summary by integrating any additional information from affect units related to the pivotal unit. By delimiting our integration to those units which are directly related to the pivotal unit, we essentially delete from our summary any information that is more peripheral to the heart of the cluster. The effectiveness of this cut-off heuristic is open to further investigation. Perhaps the cut-off boundary should be a function of cluster depth and/or span.

The actual integration of new information into the base-line summary can be handled in roughly two ways: (1) the addition of a new clause, or (2) the further refinement of existing references in the base-line summary. To see how these two techniques work, we will consider the "COMSYS" and "Broken Engagement" stories.

Both of these stories have retaliation as their pivotal units.



We will use a general frame for retaliation in our base-line summaries:

"Because Y's [??] caused a [XX] for X, X (later) [OO]ed to cause a [XX] for Y."

This frame allows us to build base-line summaries for both stories:

"Because John prevented Bill from getting a job at IBM, Bill later prevented John from getting a job."

"Because John did something bad to Mary's father, she prevented his engagement."

Notice that both of these summaries already convey the affect units for John's failure, Bill's failure, and John's loss. This is because the specification of a negative event that is part of a failure or loss unit will automatically communicate the notion of that failure or loss.

The baseline summary for the "COMSYS" story must now be augmented by the units for competition, success born of failure, and a denied request. Competition and the denied request will be integrated by a further specification of existing references. Success born of failure will require a new clause:

S:

"Because John prevented Bill from getting a job at IBM, Bill later prevented John from getting a job."

S + competition:

"Because John was promoted over Bill at IBM, Bill later prevented John from getting a job."

S + competition + success born of failure:

"Because John was promoted over Bill at IBM, Bill started his own company, and later prevented John from getting a job."

S + competition + success born of failure + denied request:

"Because John was promoted over Bill at IBM, Bill started his own company, and later refused to give John a job when he asked for one."

The base-line summary for the "Broken Engagement" story must be augmented with the units for the problem resolution and the trade-off. Both of these units will be integrated by further specification of existing references.

S:

"Because John did something bad to Mary's father, she prevented his engagement."

S + problem resolution

"Because John arrested Mary's father, she prevented his engagement."

S + problem resolution + trade-off

"Because John arrested Mary's father, she called off their engagement."

While this algorithm specifies the general structure of the summarization process, there are a number of problem areas which require extensive work:

- (1) Affect unit generational frames must be specified for both simple and complex affect units.
- (2) Knowledge-specific generational frames must be designed for those concepts which lend themselves to special verbs or constructions.
- (3) A selection process for the best generational frame must be designed.
- (4) The integration of additional information into the baseline summary must be described in detail for both further specifications and clause additions.
- (5) The interplay between (3) and (4) must be studied.

In addition to these fundamental problems within the process model, we must also examine the problems of summarization for clusters of more than one pivotal unit, and stories of more than one cluster.


9.2 SUMMARIZATION OF ARBITRARY CLUSTERS

It may be the case that stories involving more than one cluster cannot be easily reduced to a one-sentence summary. Even when a story has a single cluster, but that cluster has more than one pivotal unit, it may still be difficult to derive a one-sentence summary. These hypotheses can be tested by examining a number of stories. For now, we will look at "The Gift of the Magi" to get a sense of what the difficulties are.

Recall that there were 10 top-level affect units for this story:

[HN]	nested goals	[WN]
[HM]	regrettable mistake	[WM]
[HL1]	loss of object	[WL1]
[HL2]	loss of achievement	[WL2]
[HB]	hidden blessing	[WB]

"Loss of object" refers to the sense of loss experienced in no longer having his watch (or her hair). "Loss of achievement" refers to the sense of loss experienced in giving a gift which turns out to be a mistake (they can no longer feel good about the gifts they gave). The first sense of loss occurs as soon as the gifts are exchanged: they each realize that the gift they received is inappropriate. The second sense of loss comes with the regrettable mistake: they each realize that the gift they gave is inappropriate. The connectivity graph reveals that there are two pivotal units:



HM and WM yield maximal families, and the units for nested goals (WN

and HN) and loss of objects (WL1 and HL1) are "boundary units" in the sense that they are related to units from both maximal families. Before we discuss possible algorithms, let us look at a sample summary:

A woman sold her long locks of hair so she could buy her husband a watch chain for Christmas. But when she gave him the chain, she found out that he had sold his watch so he could buy her a comb for her hair. Initially they regretted their expensive gifts, but then they realized how much love was signified in the sacrifices made.

This summary assumes the wife's point of view for the first two sentences. Because of the symmetry in this story, it is natural to infer affect units concerning the husband as information becomes available. Let us take another look at this summary, this time identifying affect units as they are conveyed. Affect units in parentheses are inferred by shifting the perspective:

A woman sold her long locks of hair so she could buy her husband a watch chain for Christmas [WN]. But when she gave him the chain, she found out that he had sold his watch [WL2,(HL1)] so he could buy her a comb for her hair [HN,WL1,(HL2)]. Initially they regretted their expensive gifts [WM,HM], but then they realized how much love was signified in the sacrifices made [WB,HB].

The boundary units seem to be conveyed first, while the two units with the largest span occur at the very end. If this story is representative, it suggests that summaries should start with units that bound maximal families, and then proceed to those units which are more isolated later. In fact, we would have an acceptable summary if we deleted the last sentence altogether.

While pivotal units are central for stories with simple clusters, clusters with more than one maximal family are organized around the boundaries between those families. This summary started off with WN, a boundary unit from F[WM]. The other boundary unit from F[WM] is HL1. Can we build a summary starting with HL1? Try this one:

When Jim received a watch chain from Della for Christmas, he explained that he had sold his watch [HL1,(WL2)]. He sold his watch so he could buy Della a comb for her hair [HN], but he didn't know that she had sold her hair [HL2,(WL1)] in order to buy him the watch chain [WN].

In the first summary HL1 follows from WL2 by inference. In the second summary WL2 follows from HL1 by inference. Similarly, in the first summary HL2 follows from WL1 by inference, and in the second summary, WL1 follows from HL2. Since these pairs of affect units are inferentially dependent on each other, we should not allow them to distract us from the actual flow of control that is at work here. To see the pattern emerge, let us identify the HL1-WL2 pair as "X" and the HL2-WL1 pair as "Y". The order of presentation for affect units in the first summary is:

WN - X - HN - Y

The order of presentation in the second summary is:

X - HN - Y - WN

These presentations differ by a simple rotation of one unit. We could change the perspectives on these two summaries to get:

HN - Y - WN - X

Y - WN - X - HN

which gives us all four rotations. There are 20 more possible arrangements, and it is possible to generate summaries that correspond to all the permutations. So any random ordering of the four boundary

units will provide a good summary. Can we get a decent summary out of anything less than these four units? Consider:

Della had sold her long locks of hair to buy her husband a watch chain [WN], and he sold his watch to buy her a comb for her hair [HN].

This summary is based on two boundary affect units of maximal connectivity.

$$o(F[WN]) = o(F[HN]) = 3 > 2 = o(F[WL1]) = o(F[HL1]).$$

Perhaps minimal summaries can always be derived from maximally connected boundary units. To answer this question and others like it we have to study the affect analyses for a number of narratives. How often do clusters arise with two or more maximal families? Can stories with multiple clusters be reduced to single sentence summaries? Does the algorithm outlined in section 9.1 work for all simple clusters? How does symmetry affect the process of summarization? These questions can only be resolved by testing proposed algorithms on a variety of narratives.

10. RECOGNIZING AFFECT UNITS

Thus far we have explained how affect units can be used to generate summaries, but we haven't explained where the affect units come from in the first place. While the intuitive notion of affect units may be attractive, this is a useless notion for a process model unless we can specify the processes which will analyze text and produce affect units as output.

The first step in the derivation of affect units is the derivation of affect states. We cannot possibly identify an affect unit unless its component states are available to us first. The identification of an affect state is actually a fairly straightforward process, if we can assume the computational power of a predictive knowledge-based story understander. Using the knowledge structures outlined in [Schank & Abelson 1977], we can recognize affect states in terms of fairly fundamental taxonomies:

MENTAL STATE (MM)

- 1) initiating an A, D, E or I - goal
- 2) missing an enabling condition
- 3) needing a goal subsumption state
- 4) suspension or absense
of a positive interpersonal theme
- 5) a plan is intended

POSITIVE EVENT (00)

- 1) achieving an A, D, E, or I - goal
- 2) obtaining a necessary enabling condition
- 3) achieving a goal subsumption state
- 4) initiating or resuming a positive interpersonal theme
- 5) intended plan succeeds
- 6) getting news about (00: 1-5)
for some person you care about
- 7) getting news about (XX: 1-7)
for some person you dislike or hate
- 8) getting news about (MM: 1-5)
for some person you dislike or hate

NEGATIVE EVENT (XX)

- 1) A, D, E, or I - goal is thwarted
- 2) P or C - goal is initiated
- 3) script interference is encountered
- 4) initiation or intensification of
a negative interpersonal theme
- 5) intended plan fails
- 6) termination of a positive interpersonal theme
- 7) losing a necessary enabling condition
- 8) getting news about (XX: 1-7)
for some person you care about
- 9) getting news about (MM: 1-5)
for some person you care about
- 10) getting news about (00: 1-5)
for some person you dislike or hate

A discussion of this terminology would take us far afield from our central concerns, but the interested reader can find ample discussion of these references in [Schank & Abelson 1977]. Recognition for these entities has been implemented in a number of knowledge based systems, including SAM [Cullingford 1978], PAM [Wilensky 1978], and BORIS [Dyer & Lehnert 1980].

Once the three primary affect states are recognized, we can implement a predictive system of demons to build specific affect units. For example, the appearance of a mental state should construct a demon that can be activated by

- 1) another mental state (a possible m, t, or e-link)
- 2) a positive event (a possible a-link)
- 3) a negative event (a possible a-link)

If another mental state is encountered, the demon should check to see if there is a subgoal relationship (m-link), mutual exclusion (t-link), or equivalence (e-link), at work. If one of these can be verified, we have identified a primitive affect unit.

As soon as a primitive affect unit is identified, demons for complex affect units are constructed. For example, a subgoal unit should predict the possible occurrence of nested subgoals, a request, a threat, a kind act, or a malicious act. In this way, a hierarchical structure of predictions can be implemented which look for successively complicated affect units as information appears to support those possibilities.

The verification of specific affect links will rely on specific instantiations behind each affect state. There is no way of knowing whether or not two mental states should be joined by a motivation link unless we can establish a subgoal relationship between them. This subgoal relationship is naturally dependent on the specific content of each mental state. But these checks do not ask for anything that a standard inference mechanism would not need to know anyway. The information needed for these verifications should already be present in the system for inference purposes that are independent of affect unit recognition. So we are essentially constructing these units as a side-effect of existing processes. We do not need to propose additional knowledge structures to build affect units. They will fall out quite naturally from other memory manipulations with a minimal amount of extra processing.

11. CONCLUSIONS

It appears that a high level analysis for narratives can be derived from configurations of three primary affect states. These configurations consist of primitive and complex affect units whose overlapping structures allow us to measure the connectivity and symmetry of character interactions. Once the affect units in a story have been properly identified, they provide us with a framework for text summarization.

Relatively little progress has been made on the problem of text summarization. Thus far, the only program to attempt it is the FRUMP system [DeJong 1979]. FRUMP analyzes UPI stories in about 50 domains and provides summaries based on a top-down extraction of relevant

information in those domains. Summaries within a single domain do not exhibit much variation, since the summary is based on an a priori set of expectations about that domain. For example, an earthquake story will be summarized in terms of (1) where it occurred, (2) what the Richter scale registered, and (3) how many people were killed or injured. All earthquake summaries will describe those three components when they are available. This style of summarization is completely top-down and driven by specific expectations. FRUMP cannot deal with unexpected information and its summaries will reflect total ignorance of anything unexpected, even if the unexpected information is critically important.

Other top-down strategies have been proposed (although not implemented) which rely on a story grammar approach to text analysis. This approach is best typified by [Rumelhart 1976]. Rumelhart points out that a number of short narratives fall into what he calls the "EPISODE" schema. The EPISODE schema about protagonist "P" consists of:

- (1) EVENT "E" CAUSES "P" TO DESIRE GOAL "G"
- (2) "P" TRIES TO GET "G" UNTIL OUTCOME "O" OCCURS

Each of the relational terms in this schema (CAUSE, DESIRE, & TRY) refer in turn to other schema which will likewise be instantiated by particular variables within a given story. The EPISODE schema provides a root node for a hierarchical tree structure that will expand to arbitrary depth as the schemata on each level are instantiated and expanded in a recursive manner.

Rumelhart uses summarization data to illustrate how various levels of detail coincide with expansions to a particular level within the tree structure. A level 0 summary is based on the root node alone. A level 1 summary is based on the first level of expansion from the root, a level 2 summary is based on the second level of expansion, and so forth.

But Rumelhart does not discuss the recognition problem for this top-down analysis of stories. Story grammars have since been criticized for being computationally naive [Black & Wilensky 1979], and therefore of dubious value in a process model of narrative text comprehension. Aside from these processing problems, there is also the standard limitation of all top-down processors: a story grammar cannot characterize input that does not conform to its expectations. Just as FRUMP cannot deal with input outside of its knowledge domains, a story grammar would be of no help when confronted with a story whose plot was not a priori anticipated by the grammar. To what extent does "The Gift of the Magi" conform to the EPISODE schema? A hierarchical story grammar simply cannot be general enough to capture a large variation of plot structures.

A top-down analysis of narratives is not feasible. There is infinite variation in the number of plots that are possible, and people can understand stories with a new plot line in spite of the fact that they haven't seen one like it before. This suggests that plot recognition must be based on bottom-up processing rather than a top-down analysis. We can attain the flexible recognition capabilities of a bottom-up analysis scheme by constructing

configurations of primitive affect units. The information needed to recognize affect states already exists in predictive knowledge-based systems [Schank & Abelson 1977] for the purposes of inference, and the extra processing needed to link these affect states together into affect units is not difficult.

Because affect states are based on information about plans, goals and themes, affect analysis will not be applicable to stories which do not contain information along these lines. Using this approach, we will not be able to handle descriptions of sunsets or burnt steaks or waking up in the morning. Of course it is not clear that people can comfortably summarize stories that center on perceptual descriptions either, so this limitation is not a cause for concern. It would not be difficult to generate a summary that said, "This is a story describing a sunset." That is probably what a human would be reduced to as well.

In this paper we have stressed the relevance of affect analysis for the task of narrative summarization. There are undoubtedly other applications to explore. For example, a high level analysis of a story is probably used as an index into long term memory. Such an index would determine when the story can be remembered and under what conditions information from the story can be accessed [Schank 1979]. The ability to recognize similarities across stories would also be affected by a high level of narrative analysis. While processes of generation and understanding differ by much more than a procedural inverse [Meehan 1976], it might be the case that stories are generated from initial affect configurations as the starting point.

While all of these other areas are potentially relevant to affect analysis, the task of narrative summarization has the advantage of being a relatively clean I/O problem. We can give stories (input) to human subjects and ask them to produce summaries (output). This data is initially valuable in the design of a summarization process, and it will eventually allow us to test resulting programs for their psychological validity. If the process model becomes sufficiently sophisticated, it would be appropriate to study individual differences across subjects. Perhaps it would even be possible to analyze a subject's summarization behavior on a few key texts, and then predict subsequent summarization behavior on completely different texts. But it would be premature to pursue these goals right now. For the moment, it will suffice to produce a system that can generate reasonable summaries for a variety of narrative texts. This paper has attempted to show how such a goal may be realized in the near future.

BIBLIOGRAPHY

- Cullingford, R. E. (1978) "Script Application: Computer Understanding of Newspaper Stories." Department of Computer Science Research Report #116. Yale University. New Haven, Conn.
- Black, J. B., and Wilensky, R. (1979) "An Evaluation of Story Grammars." Cognitive Science. Vol. 3, No. 3, pp. 213-230.
- DeJong, G. F. (1989) "Skimming Stories in Real Time: An Experiment in Integrated Understanding." Department of Computer Science Research Report #158. Yale University.
- Dyer, M. and Lehnert, W. (1980) "Memory Organization and Search Processes for Narratives." Department of Computer Science Research Report #175. Yale University. New Haven, Conn.
- Lehnert, W. G. (1978) The Process of Question Answering. Erlbaum Associates. Hillsdale, New Jersey.
- Meehan, J. R. (1976) "The Metanovel: Writing Stories by Computer." Department of Computer Science Research Report #74. Yale University. New Haven, Conn.
- Roseman, I. (1979) "Cognitive Aspects of Emotion and Emotional Behavior" 87th Annual Convention of the American Psychological Association, New York, New York.
- Rumelhart, D. E. (1976) "Understanding and Summarizing Brief Stories." CHIP Technical Report #58. University of California at San Diego. La Jolla, Calif.
- Schank, R. C. (1975). Conceptual Information Processing. North Holland, Amsterdam.
- Schank, R. C. (1979). "Reminding and Memory Organization: An Introduction to Mops." Department of Computer Science Research Report #170. Yale University. New Haven, Conn.
- Schank, R. C., and Abelson, R. P. (1977). Scripts, Plans, Goals, and Understanding. Lawrence Erlbaum Press, Hillsdale, New Jersey.
- Wilensky, R. (1978). "Understanding Goal-based Stories." Department of Computer Science Research Report #140. Yale University, New Haven CT.

